

TOPOLOGY FOR FLEXIBLE AND PRECISE SIGNAL TIMING ADJUSTMENT

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] None.

STATEMENT REGARDING FEDERALLY SPONSORED RESEARCH OR DEVELOPMENT

[0002] Not applicable.

BACKGROUND OF THE INVENTION

Field of the Invention

[0003] The present invention relates generally to computer systems with internal operational speeds and components that impose timing constraints on clock and control signals distributed about the system. More particularly, the present invention relates to a topology for a flexible system for adjusting clock and control signal path lengths and therefore signal timing adjustments.

Background of the Invention

[0004] Modern computer systems are a conglomeration of components. Each component is designed and/or programmed to perform one or more specific tasks. However, for overall operation of the computer system to be successful, those various components must communicate with one another and exchange information. That exchange of information could be of data related to the particular application running on the system, or could be instructions of the software program itself moved from mass storage devices or random access memory to the microprocessor

for execution. Regardless of the receiving and sending device in any transfer of data or programs, there must be some coordination between the devices to make sure that information is exchanged at the proper times.

[0005] Exchange of information in modern computer systems is typically completed by the exchange of a block of information in a synchronous manner. That is, modern computer system components are interconnected by a series of bus structures capable of transmitting a predefined amount of data at any one time, *e.g.*, 64 bits of information at a time. The “synchronous” descriptor means that the exchange of information between a sending device and a receiving device is done at predefined times typically based upon a host clock signal that is available for all the components of the computer system to use as a time reference.

[0006] As computer microprocessor and bus transfer speeds increase, factors such as the speed of propagation of signals within the computer begin to come into play. For purposes of illustration, consider the generic transfer of information depicted in Figure 1 from a sending device 2 to a receiving device 4 along an arbitrary bus 6. Although the propagation of electrical signals along wires and traces of printed circuit boards is extremely fast, on the order of one inch in every 200×10^{-12} seconds (200 pico-seconds (ps)), the speed is finite. Thus, in the generic system of Figure 1, the information driven to the bus 6 by the sending device 2 is not instantaneously available at the receiving device 4; rather, the availability of the data driven by the sending device 2 at the receiving device 4 is, in part, a function of the distance between those two devices. Thus, if the sending device clock 8 and the receiving device clock 10 in Figure 1 are exactly in phase (rise and fall at exactly the same time), and there is a fairly significant distance between the two devices (on the order of several inches on a computer circuit board), then the exchange of information will not occur if the sending device 2 drives on the leading edge of the send clock 8 and the receiving

identified and testing performed on the board, the particular path that represents the signal path length closest to what was needed in the system is selected jumpering, by one of several known techniques, the connections at each end of the path desired. More particularly, the connections 20 and 22 would be electrically connected to allow the clock source to propagate along the shorter path 16, if that was the desired path, and jumpers 24 and 26 would be connected if the signal path length along the longer path 18 is desired.

[0008] However, related art implementation such as that shown in Figure 3 have several problems. First, in such a design there are only two possible signal path lengths. Thus, such a system would not compensate for the situation where the optimum signal path length is somewhere between the short path 16 and the long path 18 lengths. In such a situation, system designers typically choose the shorter path length and compensated by adding capacitance. Adding capacitance, while having the effect of slowing the rise times associated with that clock, may only be used to a certain extent before clock signal degradation becomes a problem.

[0009] Secondly, if the system of Figure 3 is used, there are several dead-end paths, or stubs, that the propagating clock signal may take. For example, the clock signal propagates from the clock source 12 out to the branch point 28. If it is assumed that the short path length is selected and jumpers are placed at locations 20 and 22 to complete the short path circuit, some of the clock signal propagates toward the open jumper 24 and reflects at that location back toward the clock signal. Likewise this occurs at the other end with regard to the open jumper location 26. These reflecting waves interfere with the clock signal and cause signal degradation.

[0010] Finally, printed circuit board space on motherboards, and the like, is a premium, thus not allowing a system designer the capability of designing in several clock paths, *e.g.*, nine or more, from which to choose later on.

[0011] Thus, what is needed in the art is a flexible and precise signal timing adjustment system that gives the system designer the maximum number of possible signal path lengths without the draw backs of using an inordinate amount of circuit board space, and without the detrimental effects associated with interference of electromagnetic signals based on reflection in stub circuits.

BRIEF SUMMARY OF THE INVENTION

[0012] The preferred embodiments relate to a structure and related method for adjusting the lengths of clock and control signal paths. More particular, the preferred embodiments relate to a topology for precisely setting the length of control and clock signal paths, and relying upon propagation times of signals along those paths to make signal timing adjustments. The preferred structure comprises at least two groups or clusters of signal paths. Each of the signal paths from the two groups or clusters preferably have different lengths such that a system designer may choose a first signal path from the first group or cluster, and a second signal path from the second group or cluster, and by selecting signal paths from the first and second group of particular lengths, the system designer may therefore control the overall length of the clock or control signal path which allows the system designer to adjust the timing of that control signal. That is, for a control or clock signal for which very little time delay is required, the system designer chooses the shortest possible path through the clusters of signal paths being an adjustable signal path circuit. On the other hand, if the system designer needs to time delay or phase lag a particular control or clock signal, the designer chooses longer signal path lengths from the first and second cluster and couples them together to make a control signal path whose length is precisely adjusted to give the desired time delay.

[0013] Preferably, the clusters of signal paths are implemented on a motherboard or other printed circuit board (PCB) in any control or clock path where timing adjustments need to be made. The motherboard or PCB card preferably has the multiple clusters of signal paths designed onto the board and the system designer selects a particular signal path from each cluster of signal paths by selectively installing zero ohm resistors. That is, each of the signal paths have ends that are proximate to an electrical contact or solder pad on the motherboard or PCB card. The system designer preferably installs zero ohm resistors from the electrical contacts or solder pads to the selected signal path, and does not install or otherwise connect to the electrical pad or solder pad to the remaining signal paths. This selective installation of zero ohm resistors is preferably done on each end of each signal path and in this way the overall adjustable signal path circuit does not have any studs or open circuit paths down which electrical waves may propagate and reflect thereby causing signal degradation in the main clock or control signal.

BRIEF DESCRIPTION OF THE DRAWINGS

[0014] For a detailed description of the preferred embodiments of the invention, reference will now be made to the accompanying drawings in which:

[0015] Figure 1 shows hardware associated with a generic transfer of information from a sending device to a receiving device;

[0016] Figure 2 shows a timing diagram of the phase relationship between a send clock and a read for the generic system shown in Figure 1;

[0017] Figure 3 shows a related art system having only two non-adjustable signal paths;

[0018] Figure 4 shows a computer system of the preferred embodiment;

[0019] Figure 5 shows an adjustable signal path circuit coupling a phase locked loop device as a signal source and a memory controller as a signal destination;

[0020] Figure 6 shows an embodiment of an adjustable signal path circuit similar to that of Figure 5;

[0021] Figure 7 shows an adjustable signal path circuit comprising three groups or clusters of available signal paths; and

[0022] Figure 8 shows an embodiment of an adjustable signal path circuit having five possible signal paths within each cluster or group.

NOTATION AND NOMENCLATURE

[0023] Certain terms are used throughout the following description and claims to refer to particular system components. As one skilled in the art will appreciate, computer companies may refer to a component by different names. This document does not intend to distinguish between components that differ in name but not function. In the following discussion and in the claims, the terms “including” and “comprising” are used in an open-ended fashion, and thus should be interpreted to mean “including, but not limited to...”. Also, the term “couple” or “couples” is intended to mean either an indirect or direct electrical connection. Thus, if a first device couples to a second device, that connection may be through a direct electrical connection, or through an indirect electrical connection via other devices and connections.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

[0024] Referring now to Figure 4, computer system 200, in accordance with the preferred embodiment preferably comprises a micro-processor or CPU 50 coupled to a main memory array

52 through an integrated bridge logic device 54. As depicted in Figure 4, the bridge logic device 54 is sometimes referred to as a "North bridge," based generally upon its location within a computer system drawing. The CPU 50 preferably couples to the bridge logic 54 via a CPU bus 56, or the bridge logic 54 may be integrated into the CPU 50. The CPU 50 preferably comprises a Pentium Pentium III® microprocessor manufactured by Intel®. It should be understood, however, that other alternative types and brands of microprocessors could be employed. Further, an embodiment of computer system 100 may include multiple processors, with each processor coupled through the CPU bus 56 to the bridge logic unit 54. To increase memory capability, and memory bus bandwidth, multiple bridge logic units 54 may be used, each coupled to its own main memory array 52.

[0025] The main memory array 52 preferably couples to the bridge logic unit 54 through a memory bus 58, and the bridge logic 54 preferably includes a memory control unit 57 that controls transactions to the main memory by asserting the necessary control signals during memory accesses. The main memory array may comprise any suitable type of memory such as dynamic random access memory (DRAM), any of the various types of DRAM devices, or any memory device that may become available in the future.

[0026] The North bridge 54 bridges various buses so that data may flow from bus to bus even though these buses may have varying protocols. In the computer system of Figure 4, the North bridge 54 couples to a primary expansion bus 60, which in the preferred embodiment is a peripheral component interconnect (PCI) bus. Figure 4 also shows a PCI device 62 coupled to the primary expansion bus 60. PCI device 62 may be any suitable device such as a modem card or a network interface card (NIC). One skilled in the art will realize that multiple PCI devices may be attached to PCI bus 60, yet for clarity of the figure, only one is shown.

[0027] A preferred embodiment of computer system 200 further includes a second bridge logic device, a South bridge 64, coupled to the primary expansion bus 60. This South bridge 64 couples, or bridges, the primary expansion bus 60 to other secondary expansion buses. These other secondary expansion buses may include an industry standard architecture (ISA) bus 66, a sub-ISA (not shown), a universal serial bus (not shown), and/or any of a variety of other buses that are available or may become available in the future. In the embodiment shown in Figure 4, the South bridge 64 bridges Basic Input Output System (BIOS) Read Only Memory (ROM) 68 to the primary expansion bus 60, therefore, programs contained in the BIOS ROM 68 are accessible by the CPU 50. Also attached to the ISA bus 66 is Super Input/Output (Super I/O) controller 70, which controls many system functions, including interfacing with various input and output devices, such as keyboard 72. The Super I/O controller 68 may further interface, for example, with a system pointing device such as a mouse 34, various serial ports (not shown) and floppy drives (not shown). The Super I/O controller is often referred to as "super" because of the many I/O functions it may perform.

[0028] The BIOS ROM 68 preferably contains firmware embedded on a ROM memory chip and performs a number of low-level functions. For example, the BIOS executes the power on self test (POST) during system initialization ("boot-up"). The POST routines test various subsystems in the computer system, isolate faults and report problems to the user. The BIOS is also responsible for loading the operating system into the computer's main system memory. Further, the BIOS handles the low-level input/output transactions to the various peripheral devices such as the hard disk drive and floppy disk drives.

[0029] Also shown in Figure 4 is a host clock 76. The host clock 76 output signal preferably couples to many of the computer system 200 components, including the CPU 10 and South bridge

64 (those connections are not shown in Figure 4 for simplicity of the drawing). Figure 4 does however show the host clock outputs coupling one each to the memory controller 58 of the North bridge 54, and to the phase locked loop (PLL) 78. The host clock signal coupled to the memory controller 58 is preferably used by the memory controller 58 to write data to the memory bus 58 during movement of data from the North bridge 54 to the main memory array 52.

[0030] The phase locked loop 78 is a device that takes a single input signal, here the host clock signal from the host clock 76, and produces a plurality of output signals having the same frequency as the input signal but shifted in phase. More particularly, the PLL has a feed-back path (not shown), the length of which controls the phase relationship between the many outputs of the PLL 78 and the PLL input. The phase relationships between the PLL 36 output and its input is a function of the length of the feed-back path between the FB Out pin and the FB In pin (not shown). An example of a phase lock loop having these characteristics is a device made by Cypress Semiconductor Corporation, part number CY 2510. Although the PLL output signals may couple to many devices within the computer, Figure 4 shows two such PLL output signals coupling one each to the memory controller 58 and the main memory array 52.

[0031] Figure 5 shows an embodiment of a signal path topology for signal timing adjustments. In particular, Figure 5 shows the phase locked loop device 78 coupled to the memory controller 58 by way of an adjustable signal path circuit 80. An embodiment of the adjustable signal path circuit 80 preferably comprises three solder pads 82, 84 and 86, which may alternatively be referred to as contact pads, electrical contacts or intersection points. These pads are preferably implemented on a printed circuit board (PCB) card such as a mother board for a computer system. These pads are preferably a highly electrically conductive material such as aluminum, which may be placed at any location on a printed circuit board as desired. Indeed, these pads 82, 84 and 86 need not be placed

on the same side of the board or even on the same level of a multiple level board. These three pads 82, 84 and 86 are preferably coupled by a plurality of traces which form signal paths, preferably having varying lengths. The combination of a solder pad and zero ohm resistors electrically connecting the various selected paths to the solder pad may be referred to as a spanning circuit inasmuch as this combination of elements spans (electrically connects) selected signal paths from each of the various clusters. More particularly, and still referring to Figure 5, an embodiment comprises a series of signal paths on a printed circuit board coupling the first pad 82 to the second pad 84. These signal paths are preferably of varying lengths and coupled on each end to the pads 82 and 84 by way of zero ohm resistors 94. By selectively installing these zero ohm resistors, which are preferably low profile devices mounted by pick and place machines during the construction of the motherboards, a system designer may choose the path length, in the exemplary embodiment of Figure 5, between the PLL 78 and the memory controller 58.

[0032] Assuming for purposes of explanation that a system designer wishes to implement path 88 between pads 82 and 84, during the process of installing the various components on the motherboard or other PCB card, only resistors 94A and 94B are installed. Thus, the PLL clock signal propagates from the control signal source PLL 78 to the pad 82 (across the unnumbered resistor which is described in more detail below), up through resistor 94A across the signal path 88, down through resistor 94B and to the second pad 84. In this exemplary embodiment, none of the resistors 94C-94F would be installed and thus no electrical path would exist along signal paths 90 and 92. Moreover, the system designer also preferably makes adjustments to the overall control signal path by selectively installing resistors 94G-94L in the second cluster of signal paths shown in Figure 5. For example, the system designer may implement the signal path 96 of the second cluster so that the overall signal path comprises the trace length from the PLL to the first pad, from

the first pad along path 88 to the second pad, from the second pad 84 along path 96 to the third pad 86, and from the third pad 86 to the control signal destination, which in Figure 5 is the memory controller 58.

[0033] It is assumed, but not required, that each of the traces 88, 90 and 92 in the first cluster, and 96, 98 and 100 in the second cluster, have a different length. Thus, a system designer may choose a plurality of different signal path lengths by selectively installing resistors 94A-94L. Still referring to Figure 5, and assuming that the signal path lengths substantially as shown in Figure 5, it is seen that the shortest path for the clock signal to travel from the PLL 78 to the memory controller 58 is through the first cluster by way of signal path 90, and through the second cluster by way of signal path 98. Likewise, in the exemplary embodiment of Figure 5, the longest path length could be implemented by the clock signal traveling through the first cluster by way of signal path 88 or 92 and through the second cluster by way of signal path 96 or 100.

[0034] Although embodiments that have duplicate signal path lengths among the various clusters are within the contemplation of this invention, in the preferred embodiment each of the signal path lengths are different, thereby allowing the system designer the maximum number of possible signal path lengths with which to tune the timing signals in the computer system 200. Referring now to Figure 6 there is shown an embodiment similar to that of Figure 5 for purposes of explaining the benefits and advantages of having signal paths with varying lengths. In particular, Figure 6 shows the signal paths of the first cluster, comprising paths A, B, and C, and the signal paths of the second cluster D, E, and F. Further assume that each of these signal paths A-F have a length of arbitrary unit as indicated in Table 1.

A	0.25
B	0.5
C	0.75
D	1.0
E	1.25
F	1.50

TABLE 1

Table 1 shows that for this embodiment, each of the signal paths A-F have a length different than the others ranging from 0.25 units to 1.50 units. The units of these lengths could be any length included but not limited to inches and centimeters. The unit of length desired is a function of the amount of delay required in the particular system. An adjustable signal path circuit 80 having the signal path lengths described in Table 1 gives a total of 9 unique signal paths through the signal path circuit. In particular, Table 2 shows each unique signal path, and that signal path's length given the arbitrary units assigned in Table 1.

AD	1.25
AE	1.5
AF	1.75
BD	1.5
BE	1.75
BF	2.0
CD	1.75
CE	2.0
CF	2.25

TABLE 2

Table 2 thus shows that for the unique path through the adjustable signal path circuit 80 comprising the signal paths A and D of Figure 6 (AD in Table 2), the total length given the assigned values in Table 1 is 1.25 units. Similarly, unique path AE has a length of 1.5 units. Thus, Table 2 shows that for the two cluster system, each cluster containing three possible paths, there are nine unique signal paths that the clock signal may take. Table 2 also exemplifies that even using the path lengths given in Table 1, where each path length is different, there are still duplicate overall path lengths. In particular, Table 2 shows that path AE is equivalent in length to path BD, path AF is

equivalent in length to BE and CD, and path BF is equivalent in length to path CE. Although an embodiment of the present invention could use an adjustable signal path circuit where some of the multiple unique paths have the same length, preferably the lengths of the signal paths are selected such that no two signal paths through the adjustable signal path circuit have the same length. While there may be many possible selections for signal paths that do not give duplicate lengths, Table 3 shows an exemplary selection for the path lengths that gives an overall adjustable signal path circuit selection ranging from 1.33 units to 4.0 units.

A	1
B	2
C	3
D	.33
E	.66
F	1.0

TABLE 3

AD	1.33
AE	1.66
AF	2.0
BD	2.33
BE	2.66
BF	3.0
CD	3.33
CE	3.66
CF	4.0

TABLE 4

Thus, it is seen that the unit lengths assigned in Table 3 (which could realistically be of the units inches) gives signal path lengths of Table 4 ranging from 1.33 units to 4.0 units, with each increment mapping to approximately 1/3 of a unit length.

[0035] Preferably then a system designer, when choosing the various components to install on a motherboard or other PCB card, will know, based on manufacturers data for those components, roughly what the signal timing characteristics for each of those components needs to be. While the system designer may have this rough idea in the stage of the engineering process where

components are selected, that designer may not know until a prototype is implemented the exact timing signal relationship. Alternatively, it is possible that motherboards or PCB cards are designed based upon a component selection that may change. For example, many microprocessors may plug into a standard zero insertion force (ZIF) socket. That is, microprocessors from Intel® as well as AMD may each fit into a particular location on a motherboard or PCB card, but each may have differing signal timing requirements. Thus, a motherboard or PCB card, whose layout and design is finalized well in advance of the system designer knowing what components may be used, preferably implements an adjustable signal path circuit such as 80 with the path lengths selected such that the system designer has a range of possible lengths to implement depending on parameters and variables that are not determined at the time of the motherboard or PCB card finalization. Once the signal timing requirements are determined for a particular motherboard or PCB card configuration, the system designer chooses a particular path in the adjustable signal path length circuit to accommodate that signal timing. Thereafter, each motherboard or PCB card populated with various devices selectively has the zero ohm resistors populated only for those signal paths desired in the particular implementation.

[0036] Referring still to Figure 6, it is noted that the topology given in Figure 6, and those equivalent to it, have an advantage that for any signal path that is not implemented (that is for any signal path for which the zero ohm resistors on its path beginning and path end are not installed), there are no dead-end paths or stubs. As clock signals propagate along a wire or trace on a PCB card, the signal splits at each junction and propagates in each direction. An open circuit acts as a reflector for that signal. That is, the portion of the clock signal that propagates along the dead-end path, reflects at the end and propagates back toward the junction. When the reflective signal meets with the continuing clock signal, interference occurs which may degrade signal integrity in the

rather than inches). Using the signal path lengths given in Table 5 for the three cluster arrangement, the system designer has 27 possible signal path lengths ranging from 11.33 units to 20.0 units in 0.33 unit steps.

[0039] Thus, it is seen that the embodiments of the invention give the system designer a topology for signal timing adjustments that is highly flexible and may be easily and precisely tuned for the particular components on the motherboard or PCB card. Referring back to the computer system 200 shown in Figure 4, an adjustable signal path circuit 80 could be implemented anywhere in the computer system where the system designer needs to vary the length of a signal path, be it for a clock circuit or any other control signal propagating within the computer system. Preferably, however, the timing constraints between the memory controller 58 and the main memory array 52 may require that the clock signals feeding each of these devices (it is noted that the memory controller 58 has a clock signal both from the host clock 76 and the PLL 78) preferably implement one of these variable length signal path circuits so that the timing signals for reads and writes between them may be adjusted.

[0040] All the various embodiments shown in the drawings (Figures 5-8) show zero ohm resistors on the path that leads to the first pad of the first cluster, and zero ohm resistors coupling the trace leading away the last pad of the last cluster. More particularly, and referring to Figure 5, the adjustable signal path circuit 80 is shown to have a zero ohm resistor 102 coupled between the PLL 78 and the first pad 82 of the first cluster of signal paths. Likewise, the adjustable signal path circuit 80 also has another zero ohm resistor coupling the last pad 86 of the second cluster to the memory controller 58. In the specific embodiment shown in Figure 5, the resistors 102 and 104 may be replaced simply by coupling the PLL 78 to the pad 82 directly, and likewise coupling the pad 86 to the memory controller 58 directly. However, these resistors are shown in Figure 5 (and

the remaining embodiments shown in Figures 6-8) to exemplify that it is not necessary that only a single adjustable signal path circuit be used in any particular location. That is, if the system designer is unsure of the length required for the signal path that may be accommodated by one of the adjustable signal path circuits 80, it is possible that multiple adjustable signal path circuits 80 may be placed in parallel on a motherboard or PCB card. As components are selected and signal timing relationships solidify into more distinct ranges, the system designer may implement any one of these parallel implementations of adjustable signal path circuits by insuring, at the time of population of the motherboard or PCB card, that the zero ohm resistors coupling the traces are installed on the adjustable signal path circuit which provides a range closest to the calculated or estimated signal path length.

[0041] As discussed with respect to the embodiments exemplified in Figures 5-7, each cluster of signal paths in these adjustable signal path circuits 80 have three possible signal pads for the clock or other control signals to travel across each cluster. In embodiment of Figure 7, it is shown that an additional cluster, in Figure 7 to make a total of three clusters, may be added to increase the total number of possible selections from the adjustable signal path circuit. However, it must be understood that the embodiments of the present invention are not limited to signal path circuits having clusters with only three signal paths. Figure 8 shows another two cluster embodiment of an adjustable signal path circuit 80 where each cluster has a total of five possible signal paths within each cluster. A two cluster, five signal path per cluster embodiment, such as that shown in Figure 8, gives 25 possible unique paths. Just like the embodiments where additional clusters are added to the three signal path clusters circuits, additional clusters may be added to this five signal path cluster embodiment for additional flexibility in the design of signal and clock path lengths.

